

Procedures and Software for Assessing Uncertainty in Cost Estimates

Henry L. Eskew • Walter R. Nunn


19990429 034

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Center for Naval Analyses

4401 Ford Avenue • Alexandria, Virginia 22302-1498

Approved for distribution:



Paul E. Speer
Director, Resources Team
Support Planning and Management Division

This document represents the best opinion of CNA at the time of issue.
It does not necessarily represent the opinion of the Department of the Navy.

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

For copies of this document, call the CNA Document Control and Distribution Section (703) 8

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 1995	3. REPORT TYPE AND DATES COVERED Final	
4. TITLE AND SUBTITLE Procedures and Software for Assessing Uncertainty in Cost Estimates			5. FUNDING NUMBERS C - N00014-91-C-0002	
6. AUTHOR(S) HL Eskew, WR Nunn				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Center for Naval Analyses 4401 Ford Avenue Alexandria, Virginia 22302-1498			8. PERFORMING ORGANIZATION REPORT NUMBER CRM 95-87	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 Words) With extensive literature on the treatment of uncertainty in cost estimates, and because several uncertainty software packages are now available, the Naval Center for Cost Analysis asked the Center for Naval Analyses to conduct a study with the objective of evaluating the procedures and software that it now employs. This report begins with a discussion of introductory analytical issues and then focuses on the software packages that were evaluated. Additional analytical questions are addressed in connection with those evaluations.				
14. SUBJECT TERMS assessments, computer programs, cost estimates, costs, defense economics, economic theory, economics			15. NUMBER OF PAGES 46	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Contents

Introduction and summary	1
Analytical preliminaries	5
A note on concept and terminology	5
A simple example of cost uncertainty analysis	6
RISK	11
Workscreens and further analytical issues	13
Crystal Ball and NCAP	23
Detailed comparison of Crystal Ball and NCAP	23
Documentation	24
Running time	24
Size limitations	24
Number of variables analyzed	24
Ability to handle correlations	27
Distributions of random variables	28
Concluding remarks	28
Appendix A: Literature review	29
Introduction	29
Discussion	30
Definition of "risk"	30
No methodology is "best".	31
Probability distributions	32
Correlations between costs	33
Data collection and analysis	33
Comparison of selected software packages	33
CLT—Central Limit Theorem (USASSDC) [7].	33
FRISKEM—Formal Risk Evaluation Methodology (Aerospace) [18]	34
PACER—Parametric Cost-Estimating Relationship Module (DSMC) [7]	34

@RISK—Spreadsheet Add-in Model (Palisade Corp.) [7]	35
Crystal Ball—Spreadsheet Add-in Model (Decisioneering, Inc.) [4, 19]	35
RISK—Cost Risk Model (Tecalote Research) [7].	35
Appendix B: Parameter computations involving the product of two random variables	37
References	41
Distribution list	43

Introduction and summary

The Naval Center for Cost Analysis (NCA) is a major contributor to the cost analyses of all major Department of the Navy acquisition programs. Those analyses are submitted to the OSD Cost Analysis Improvement Group (CAIG) for review. The quality and defensibility of the estimates are crucial to a program's receipt of approval to proceed beyond each acquisition milestone.

Present Department of Defense (DOD) directives require that the uncertainty associated with the cost estimates be quantified and displayed for the CAIG review. Accordingly, NCA has developed and is now using a set of statistical procedures, which have been embedded into an electronic spreadsheet package, for assessing uncertainty in the estimates. Nevertheless, because there is an extensive literature on the treatment of uncertainty in cost estimates, and because several uncertainty software packages are now available—both from commercial sources and from various DOD organizations—NCA asked CNA to conduct a study with the objective of evaluating the procedures and software that it now employs. This is the final report of that study.

We began with a brief literature search aimed at (1) obtaining a perspective on the state of the art in this area, and (2) becoming acquainted with the capabilities of the software packages that appear to be in common use. Appendix A contains a report of the search. We followed that with a briefing to NCA. At the conclusion of the briefing, it was mutually agreed that we would explore two analytical issues bearing on the assessment of cost uncertainty, and that we would further evaluate a specified subset of the packages. The subset included the package presently in use at NCA.

The analytical issues were (1) the nature and proper treatment of correlation among cost elements, and (2) the types of probability distributions that best characterize uncertainty in cost estimates under

different circumstances. A third issue, choosing measures of variability (dispersion) for use in the distributions, arose in the course of the work and is also addressed in this document. The packages selected were:

- **RI\$K**—This software was developed by Tecolote Research, Inc., for inclusion in the Automated Cost Estimator Integrated Tools (ACE-IT), a system developed by joint funding from the Army and Air Force. RI\$K also operates in a standalone mode.
- **Crystal Ball**—This commercial package was developed by, and is licensed by, Decisioneering, Inc.
- **NCAP**—We use this title to refer to the package developed for use within NCA by Richard L. Coleman, Captain, USN (Ret.).

A second commercial package, @RISK, was originally included in the subset, but after further review of its documentation and conversations with its developer, Palisade Corp., we decided that because of its close similarity to Crystal Ball, evaluation of only the latter would meet the needs of the study.

Our principal findings and conclusions may be summarized as follows:

- Correlations between cost elements are important; they should not be ignored. When the source of the correlation is a direct linkage between a driver cost and a dependent one, the correlation can be adequately reflected by any of several methods. When the linkages between two or more costs cannot be made explicit, or whether such linkages exist at all, is a matter of considerable controversy. Some analysts reject outright the use of subjective measures of correlation; others strongly encourage it. A middle ground is that sensitivity analysis can inform the debate in any particular case.
- When cost estimates are generated by linear (log-linear) regression equations, and the standard assumption is made that the error term is normal (lognormal), we believe the normal (lognormal) distribution to be appropriate forms for characterizing the uncertainty associated with the estimates, given that

certain adjustments are made relating to the t distribution. In those same cases, we also believe that the prediction error is the correct measure of variability (dispersion) because it incorporates all sources of uncertainty inherent in the regression. In many other cases, where the estimates are generated by methods other than regression, both argument and evidence support the use of right-skewed distributions (e.g., lognormal, triangular, or beta).

- The RI\$K software package, which requires no electronic spreadsheet, has many attractive features and is continually being improved. Its user's guide also provides a thorough tutorial on cost uncertainty analysis. NCA has immediate access to RI\$K, and we think analysts can profit from its use and its documentation.
- Crystal Ball is our preference as a spreadsheet overlay. It is well documented, powerful, flexible, and easy to use, and it facilitates documentation of an uncertainty analysis. It is relatively inexpensive, but nonetheless it must be purchased.

The report begins with a discussion of introductory analytical issues. We then focus on the software packages that we evaluated. Additional analytical questions are addressed in connection with those evaluations.

Analytical preliminaries

A note on concept and terminology

Before proceeding with a discussion of analytics, we consider it important to elaborate on the concept of cost uncertainty and to comment on a question of terminology. Virtually without exception, cost estimates take the form of point estimates—"Our estimate is that the cost of Engineering and Manufacturing Development (EMD) will be \$325.4 million (FY 1995 dollars)." Unfortunately, the one thing known with complete certainty about such a statement is that it will prove to be wrong. Cost estimation is in no sense an exact science. A far more realistic and useful perspective is to think of the point estimate as simply one outcome in a range of possible outcomes. Many factors contribute to the width of such a range, and to the relative likelihood that the final outcome (cost) will fall within various portions of the range. The task of cost uncertainty analysis is to quantify those ranges and relative likelihoods. In short, it becomes an exercise in the application of probability theory and methods to (1) empirical data bases, and (2) information specific to the program for which the estimates are being developed.

The issue in terminology has to do with the distinction between *risk* and *uncertainty*. What makes this an issue is that those terms are used inconsistently—and sometimes interchangeably—in the professional literature, in documentation accompanying software packages, and in various government publications. At the expense of some oversimplification, there appear to be three positions on the matter. The first is that a program's costs are influenced by several (perhaps very different) sources of uncertainty, and the process of quantifying the effects of those influences through probabilistic modeling is called risk analysis. A second view is that risk has to do with the cost impact of potential variability in a program's schedule or its design and technical characteristics, whereas uncertainty arises from inherent

limitations in the data and methods available to the cost analyst. A third interpretation is that the two terms are synonymous. None of these positions seem unreasonable to us. Because NCA subscribes in general to the second, we have chosen to do likewise. We do note, however, that most of the available software packages that support this kind of analysis have the term *risk* in their titles.

A simple example of cost uncertainty analysis

To set the stage for the subsequent discussion, we provide the following highly simplified example of cost uncertainty analysis. One purpose of the example is to highlight the role of probability distributions and Monte Carlo simulations, as well as the effects of interdependence (correlation) among cost elements. Another is to lay the statistical groundwork for the discussion that follows.

Consider a work breakdown structure (WBS) that consists of only two cost elements: hardware (*H*) and support (*S*). The sum of the two equals total cost (*TC*). Suppose we had reason to believe that the uncertainty associated with each element could be characterized by normal probability distributions having parameter values as follows:

Table 1. Hypothetical parameter values

Element	Mean (μ)	Std. deviation (σ)
Hardware	100	20
Support	50	10

We are ultimately interested in the parameter values and distribution of total cost. From the definition of sums of random variables, we know the following:

$$\text{Mean (TC)} = \mu_{TC} = \mu_H + \mu_S = 100 + 50 = 150$$

$$\text{Standard deviation (TC)} = \sigma_{TC} = (\sigma_H^2 + \sigma_S^2 + 2\rho\sigma_H\sigma_S)^{1/2} = (400 + 100 + 400\rho)^{1/2},$$

where ρ is the correlation between hardware and support, and $\rho\sigma_H\sigma_S$ is the covariance between *H* and *S*. Because by definition, $-1 \leq \rho \leq 1$, that parameter has a very important influence on the size of σ_{TC} and

thus on the uncertainty associated with total cost. At the extremes, σ_{TC} could be as low as 10 or as high as 30. We are therefore unable to proceed with the uncertainty analysis without dealing in some fashion with the correlation between the two cost elements. In actual practice, of course, the treatment of correlation between any two cost elements would depend on both the nature of the elements and the particulars of the program for which the estimates are being developed. For the expository purposes of this section, we consider four possibilities. Later in the paper, we discuss another two.

The simplest thing to do is to assume that the two elements vary independently, i.e., $\rho = 0$. Because components of support costs frequently have direct ties to hardware costs, that assumption hardly seems plausible in this example. In general, however, it may be quite reasonable to posit that two or more cost elements are uncorrelated. Maintaining for a moment the assumption of independence, there are two ways of proceeding from this point. One is the Monte Carlo approach. A fairly large number (1,000 or more) of random drawings would be taken from the postulated hardware and support distributions, and the two sets would be added—starting with the first pair of drawings and ending with the last—to form the distribution of total cost. The mean, standard deviation, and percentiles of the cumulative distribution would be computed, making possible statements such as, “We’re 90 percent confident that total cost will not exceed \$410 million.” This would then complete the uncertainty analysis. Alternatively, because H and S are both normally distributed, we would be very safe in assuming TC to also be normal. We could compute the mean and standard deviation of that distribution as shown above, and by referring to a table of standard normal values, calculate percentiles without resorting to simulation. This approach is typically called *analytic* or *heuristic*. When there are several different forms and shapes of probability distributions involved in an uncertainty analysis, Monte Carlo simulation is generally thought to be preferable. Nevertheless, the alternative approach is much simpler to execute and in many cases provides results that are extremely close to those generated by the simulation.¹

-
1. The authors of [1] describe experimental evidence showing that heuristic methods, with total cost assumed to be normal, provide excellent approximations to the simulated distributions. Those results are fairly robust across numbers of cost elements, degrees of skewness in the cost element distributions, and degrees of correlation among elements.

A second possibility is that support costs are being estimated as a fixed fraction (factor) of hardware costs. (The numbers in table 1 are consistent with a factor value of 0.5.) In that case, with S being simply a linear transformation of H , ρ is identically equal to 1.0. As with the preceding case, the total cost distribution could be obtained either by simulation or heuristically. For the former, a large number of random drawings would be taken from the hardware distribution, and then each value in the set would be multiplied by 0.5 to obtain the distribution of S . The two sets would be added as before to generate the distribution of total cost. Correlation of 1.0 between H and S is therefore built into those two variables. For the heuristic approach, parameters of TC could be calculated directly and the remainder of the process carried out as described above.

A third case is where support costs are estimated as a fraction of hardware costs, but there is uncertainty as to the magnitude of the factor. What is typically done in such cases is to treat the factor as a random variable, and to specify a distribution form and parameter values for it. The consequences are that:

- The standard deviation of support costs will increase from its previous value because S now reflects the combined variability of H and the factor, and
- The correlation between H and S will decline from its previous value of 1.0 because the interdependence of the two is no longer exact.

It is possible to deduce analytically the new value of σ_s and the value of the correlation coefficient. Those computations are shown in appendix B. The two parameters may also be obtained by simulation. In table 2, where results of these three approaches are compared, the variable factor in the third case was assumed to be uniformly distributed over the interval $[0.35, 0.65]$. All results in the table were obtained analytically.

Table 2. Comparison of alternative approaches

Support costs	Std. deviation - support	Correlation coefficient	Std. deviation - total cost	Total cost at 90% confidence
Independent of hardware	10.0	0.00	22.4	179
Fixed fraction of hardware	10.0	1.00	30.0	188
Random fraction of hardware: $U(0.35, 0.65)$	13.3	0.75	31.3	190

The fourth case leads into what is probably the most controversial area of cost uncertainty analysis. In the context of the present example, the situation would be that hardware and support costs cannot be linked by any factor relationship or other explicit mechanism, but they nevertheless are believed to move together—to be correlated. The underlying source of the correlation, while maybe not totally obscure, simply does not lend itself to incorporation in a set of cost-estimating equations. Examples that appear in the literature, and which apply to different phases of life-cycle cost, have to do with slipped schedules; failure to achieve technical breakthroughs; unforeseen business-base conditions; and policy changes affecting deployment, operations, and logistics support. Some analysts find this totally reasonable and are quite ready to provide subjective measures, if necessary, of the degree of interrelatedness among cost elements. Those analysts necessarily require that their supporting software makes provision for introducing correlation in this fashion. Other analysts take one or the other of the following positions, or possibly a combination of both:

- Subjective estimates of correlations have no place in a cost uncertainty analysis. If an interrelationship exists and has not been made explicit, the cost model is deficient.
- Schedules, technical breakthroughs, etc., constitute *risk*, not cost *uncertainty*, and they should be dealt with in a separate analysis. Subjective estimates may be used in the separate risk analysis.

If forced to take a side in the controversy, we would probably side with the subjectivists for four reasons. First, the basic argument is compelling. Second, a great deal of any uncertainty analysis involves subjective judgment; there is nothing unique about subjective quantifications of correlations. Third, software packages that permit explicit introduction of correlation coefficients are more flexible than those that do not. And finally, there is always the possibility of conducting sensitivity analyses on the correlations. Such analyses may reveal in any given situation that the issue is moot.

We turn now to the three packages that we evaluated: RISK, NCAP, and Crystal Ball, and to certain additional analytical issues.

RI\$K

As mentioned earlier, RI\$K is available as a tool in ACE-IT and can also operate in a standalone mode. The following is a summary of what we consider to be the principal features and strengths of RI\$K:

- *Development.* Unlike commercial packages such as Crystal Ball, which are designed for application in any field of science and engineering, RI\$K was developed by a group of experienced cost analysts and statisticians for use by other cost analysts. It includes various options and defaults based on analysis of empirical cost and programmatic data.
- *Documentation.* The user's guide accompanying RI\$K accomplishes two objectives: (1) it makes the software easy to use, and (2) it serves as a thorough tutorial on conducting cost uncertainty (risk) analysis.
- *Electronic spreadsheet.* Many of the packages we reviewed can operate only as overlays to the standard spreadsheets, e.g., Lotus 1-2-3 or Excel. RI\$K, on the other hand, is self-contained; it comes with what is essentially its own spreadsheet.²
- *Probability distributions.* RI\$K accommodates five forms of probability distributions: normal, lognormal, triangular, beta, and uniform. Something of an ad hoc procedure can be applied in using the normal distribution when a *t* distribution is technically more appropriate.
- *Correlation.* Provision is made, although with some limitations, for explicit introduction of measures of correlation among cost elements. Users of the package are encouraged, however, to

2. A further explanation of this statement is that each of the RI\$K work-screens, which are discussed later, is in fact a subset of columns from a single spreadsheet.

think of (and to formulate) such measures as subjective indicators of the strength and direction of "group associations" rather than as strict product-moment correlation coefficients.

- *Method.* RI\$K offers a choice between Monte Carlo simulation and a closed-form analytic method for generating aggregate distributions. If the latter is selected, the software assumes that the desired distribution can be adequately described by a beta curve. If there is interest in the extreme tails of the distribution, the Monte Carlo method is recommended—with a very large number of random drawings. (A user can specify the number of drawings desired.)
- *Output.* Tabular output includes, for each WBS element, basic statistics (means, medians, standard deviations, etc.), confidence (percentile) levels, correlation tables, and user inputs. In addition, histograms of the derived probability distributions and continuous graphs of cumulative distributions are available for any WBS elements desired.

The preceding features were characterized as strengths of the RI\$K software. There are certain other features that, at least in our opinion, constitute weaknesses. Before describing those, we should note that RI\$K, like its parent system, is continually evolving. Discussions with its support contractor, Tecolote Research, Inc., revealed that work is either under way or could easily be carried out to remedy the major weaknesses.

Although RI\$K permits one cost element to be estimated as a fraction (factor) of another, and although a probability distribution can be placed on the factor (with moderate restrictions), the software is not designed for a user to introduce more complicated linkage equations. For analysts who subscribe to the philosophy that the only legitimate correlations are those that arise from explicit linkages of cost elements, this is a near fatal flaw in RI\$K. Our own view, which we will justify later in the paper, is that with the proper choice of variability (dispersion) measures, and by occasionally resorting to off-line simulation or "tricking" the software into accepting a more complicated equation, this weakness in the model's present configuration can be

overcome. And as suggested above, effort is under way to incorporate more general solutions.

Other than the limitations imposed by the absence of a t distribution and by the restrictions placed on factor distributions, the remaining feature that we find inhibiting is the way RI\$K handles correlations or "group associations." Imagine that WBS element E_1 drives elements E_2 and E_3 , but at different strengths of association. Although not highly likely, it is entirely conceivable that an analyst knows the three pair-wise correlations and wishes to provide them as input to the analysis. RI\$K prohibits that in effect by requiring, for purposes of supplying correlation measures, that each group of elements be mutually exclusive of every other group. Thus if E_2 and E_3 were in a group with E_1 , their own correlation could not be specified in a second group, nor could they be in a group with other elements. We are quite prepared to believe that this could make little or no difference in many real uncertainty analyses, but it is nonetheless a limitation that is not encountered in other packages such as Crystal Ball.

Workscreens and further analytical issues

RI\$K is structured around a series of five workscreens. As noted in table 3, which provides an overview of the screens and the role played by each, only two of the five are absolutely essential in every analysis. We will provide further observations on the workscreens, with some of those constituting the springboard for discussion of additional and important analytical issues.

As for the Initial Estimate screen, another convenient feature of RI\$K is that the WBS hierarchy—the successive levels of aggregation—is defined simply by each element's order of entry and level of indentation. There is no requirement to write summation expressions nor to document the location (in a spreadsheet, for instance) of any variables. Concerning documentation in general, the combination of printed copies of the workscreens and the various forms of output constitute complete documentation of the uncertainty analysis—at least from the point of view of reproducibility.³ An analyst would

3. Presumably, the baseline cost estimate would be documented elsewhere.

probably want to provide additional documentation on the choice of probability distributions and the origins of measures of dispersion and group associations (if applicable).

Table 3. RI\$K workscreens

Workscreen	Required or optional	Inputs	Comments
Initial estimate	Required	WBS, baseline cost estimate, and types of estimation method used	Sequence of entry and level of indenture convey WBS hierarchy
Estimating risk	Required	Forms of probability distributions and measures of dispersion and skewness	Inputs are supplied for each element that is not the sum of a set of subordinate elements
Other risk	Optional	Characterization of schedule, technical, and configuration risk	All inputs are subjective
Factor specifications	Optional (req'd if using factors)	Identification of driver elements for those costs estimated by factors	Probability distributions may be specified for factors
Groupings	Optional	Identification of group associations and strengths among WBS elements	One element may be designated as dominant in a group

Inputs to the Estimating Risk screen specify the type of probability distribution chosen for each element that requires one, together with measures of dispersion and skewness. For some types of distributions, the dispersion measures are quantitative; for others, they are subjective. All measures of skewness are subjective. These observations prompt the following discussion of selecting distribution forms and measures of dispersion.

With regard to choosing distribution forms, a few things seem relatively clear. Statistical regression analysis plays a central role in (1) developing baseline cost estimates, and (2) providing a basis for quantifying the uncertainty associated with the estimates. If a strictly linear regression equation serves as the mechanism for estimating a

particular cost, conditional on a set of values for the equation's predictor or driver variables, then the uncertainty associated with that estimate (prediction) should ideally be characterized by a t distribution. If that distribution is not available in the package, the normal distribution should be used with adjustments as noted in a subsequent paragraph.⁴ The reason for this choice is that the regression equation arises from a model that assumes the presence of a normally distributed random error term. (Because the variance of the error term is unknown and must be estimated, the relevant distribution, including the distribution of the prediction, becomes the t rather than the normal.) Regression equations that are linear in the logarithms of their variables are also widely used in cost analysis. There the underlying assumption is that the error term is lognormally distributed, and for uncertainty purposes, the lognormal seems the correct choice—again with certain adjustments that pertain to the t distribution.⁵

There are two other attractive features of the lognormal, whether in connection with a log-linear regression or as a characterization of uncertainty for an estimate developed by some other method. One

4. The importance of the distinction between t and normal distributions is one of sample size or, more precisely, degrees of freedom (d.f.), in the database from which the regression equation was developed. At 30 d.f., the t and the normal are essentially equivalent at one decimal place. As the number of degrees of freedom become quite large, the t converges to the normal.
5. When a lognormal distribution is specified, RISK interprets the baseline cost as a *median*. Because of the distribution's right-skewness, the median is lower than the mean. The package employs a procedure for increasing the baseline to a mean cost for the element, because the sum of the means of the cost elements is what constitutes the mean of total cost. We have only a minor objection to this, in the case where a lognormal is specified because a log-linear regression is being used. The baseline estimate produced by the regression, while appearing to be a median, is in fact an upwardly biased estimate of the median. See [2], appendix A. The magnitude of the bias could be small or large, depending on a variety of factors. It may well be that interpreting the baseline as a mean would prove more accurate in the long run, but there is certainly no way of demonstrating that.

feature is its right-skewness and the other is that the lognormal precludes a cost variable from becoming negative. There is a wealth of experience indicating that when costs are under-predicted, the magnitude of the errors is considerably greater than when they are over-predicted. Right-skewness in a distribution provides a means of capturing that phenomenon. One also finds evidence of right-skewness in [3] with respect to costs that are usually estimated by factor relationships. Examples are engineering change orders and initial spares.

A situation that involves both right-skewness and negative values is one in which, for whatever reason, the measure of dispersion is quite large relative to the baseline cost estimate. By "quite large" we mean greater than 50 percent. Choice of a normal distribution in such cases seems particularly unwise because some nontrivial fraction of the cost values would be negative. Either a lognormal or a right-skewed beta or triangular distribution might be more sensible.⁶ Appendix A references the favorable discussion of beta and triangular distributions in the literature, but also mentions the problems associated with accurately specifying the finite upper and lower values of those distributions. We note that RISK requires, as input, that the "spread" of the triangular and beta be described only as low, medium, or high, and that skewness be described simply as right, left, or center. The package's numerical default values are documented in its user's guide.


We turn now to the issue of choosing measures of dispersion. Our discussion is confined to those cases where regression equations are used to generate the baseline cost estimate. Some explanation of figure 1, where the measure is depicted, will facilitate what follows.

The figure assumes the existence of a database from which a simple linear regression equation—the upwardly sloping straight line—has been developed. The cost driver is X , having a mean of \bar{X} in the sample. The value of the driver for purposes of prediction is X_0 . The *point* estimate, or prediction, is the point on the regression line corresponding to X_0 . The hyperbolic curves represent the width of the prediction interval at a specified level of confidence. The interval is

6. When either a triangular or beta distribution is specified, RISK interprets the baseline cost as a *modal* value.

smallest at the sample mean of the driver, and becomes progressively larger as X_0 moves away from the mean in either direction. It is computed by first multiplying the prediction error (PE) by the appropriate value of t , given the number of degrees of freedom and the desired confidence level, and then adding that result to, and subtracting it from, the point estimate. (We note that the RISK user's guide refers to what we call the prediction error as the prediction *interval*. We consider that an unfortunate choice of terminology because the prediction interval is defined as we define it here in all statistical and econometric literature with which we are familiar.)

Figure 1. Prediction intervals and prediction errors



The PE captures all sources of uncertainty embedded in the prediction, except for any uncertainty associated with the value of X_0 . Those sources are:

- Variance of the estimate of the intercept parameter
- Variance of the estimate of the slope parameter
- Covariance between the intercept and slope estimates
- Variance of the model's random error term.

If RI\$K (or any other package) provided a t distribution, the PE would be the correct measure of dispersion for the baseline (predicted) cost. It is directly analogous to a standard deviation. When only a normal distribution is available, the PE must be adjusted upward to reflect the fact that the t distribution has thicker tails than the normal. An example of such an adjustment is provided later in the discussion.

Two qualifying remarks are in order with regard to the preceding few paragraphs. First, if the regression is log-linear, the same process applies. A specialized feature of RI\$K, however, is that when a lognormal distribution is specified, the baseline cost is interpreted as being in dollars, but the dispersion measure is expected to be in decimal form, i.e., unchanged from the value that was generated in log space. Second, even in the case of a single driver variable in a regression equation, and certainly with multiple drivers, the analyst may not have all the information needed to compute prediction errors for input to the analysis.⁷ This is not a serious problem with a single driver, and the RI\$K user's guide provides a table of approximate adjustment factors, but it can definitely be a problem with multiple predictor variables. The documentation associated with virtually any regression equation will include the S.E.E. About the best that can be done is to make a subjective upward adjustment to that value, taking into consideration the degrees of freedom and the extent to which the variables are perceived (if not actually known) to deviate from their sample means.

We've made several references to, and provided examples of, simple factor relationships where one cost variable drives another. In

7. This limitation was first noted by Vern Reisenleiter of NCA.

addition, a relationship between two cost variables frequently arises from a linear regression analysis. Thus the prediction of the dependent cost (C_y), rather than being conditioned on a given X_0 , becomes a function of a driver (C_x) that is itself subject to uncertainty. The approach we recommend for simulating C_y is as follows. Letting b_0 and b_1 represent the estimates of the intercept and slope parameters, respectively, and PE_a the adjusted prediction error, the analyst should form the equation

$$C_y = b_0 + b_1 C_x + E ,$$

where E is the random variable capturing the uncertainty associated with the regression. Then take random drawings from the distribution of C_x and multiply each by b_1 . Continue by taking random drawings from E , whose mean is zero and measure of dispersion PE_a . Each random drawing of E is added to the corresponding value of $b_1 C_x$, along with b_0 . C_y will have the same mean (except for sampling error) as it would if C_x were a constant, and its variability will reflect the uncertainty from the regression and the uncertainty associated with C_x . Although RI\$K is not designed to accept equations such as this, it can be "tricked" into doing so by setting up three subelements of cost under C_y , each of which represents one component of the above equation.⁸ Of course, the simulation could be done off line, and the resultant standard deviation of C_y and correlation between C_x and C_y could be input directly into RI\$K without further complication.

The following example will tie much of the preceding discussion together. It is taken from a recent NCA cost analysis. The dependent cost is engineering design, and the driver is total EMD hardware cost. The regression equation, in millions of FY 1993 dollars, was

$$C_y = -0.016 + 0.84 C_x .$$

-
8. The trickery turns out to be a little more complicated than described above. With the mean of E equal to zero, half of its values will be negative. Because RI\$K truncates negative values in a distribution, the user receives an error message saying that too many values are being truncated. This can be overcome by arbitrarily choosing a mean of E that is high enough to avoid negative values, and then subtracting that mean from the constant b_0 . RI\$K will accept negative values if they are constants.

The standard error of estimate (S.E.E.) was 8.8, and there were 3 degrees of freedom. To obtain the prediction error, we multiply the S.E.E. by the value of the square-root term shown in figure 1 (1.13 in this case). The result was 9.9. We then adjust that upward in order to shift from the t to the normal distribution. The adjustment factor is unique to a given confidence level and degrees of freedom. For illustration here, we chose 90-percent cumulative confidence, meaning 10 percent of the distribution is to the right of that point. The factor is then computed as the ratio of t_{10} for 3 d.f. (1.638), to the value of a standard normal variate with the same area to the right (1.282). The resulting ratio was 1.28. Thus the adjusted PE was $1.28(9.9) = 12.7$.

With C_x held constant at its mean of 28.0, the baseline estimate for C_y was

$$C_y = -0.016 + 0.84(28.0) = 23.5 .$$

The 90-percent cumulative confidence for that cost, determined analytically, was 39.8. The computation is

$$23.5 + 1.282(12.7) = 39.8 .$$

This reflects only the uncertainty inherent in the regression. (Note that this number is nearly 70 percent larger than the baseline estimate. Considerable uncertainty is associated with this cost, owing to a relatively large S.E.E. and a very small sample size.) When C_x is allowed to vary randomly in accordance with its (normal) distribution, the baseline value of C_y is unaffected, and the combined measure of dispersion for C_y (reflecting both regression and driver uncertainty) increased to only 13.0, a value obtained by simulation. The corresponding cost at 90-percent confidence was 40.2, just slightly larger than before. The closeness of the two sets of results is attributable to the minimal uncertainty (low dispersion measure) associated with C_x . That, of course, is simply a feature of the particular example we chose, and it definitely need not be the case in general.

Returning to the discussion of workscreens, the Other Risk screen provides the user a means for incorporating additional uncertainty if the program in question is thought to face schedule and technical requirements that are unusually difficult in comparison with

programs of similar types. As noted in table 3, all inputs are subjective. These increases in uncertainty may be thought of as *penalty factors*. RI\$K provides default penalty factors for two system types, hardware and software.


The Factor Specification workscreen is used to designate one cost element as a simple factor of another.⁹ Naturally, if there are no factor relationships, this screen is not required. When the estimation method for an element is specified as *factor* in the Initial Estimate screen, and when that element is linked to a driver in the Factor Specification screen, RI\$K calculates the value of the factor by computing the ratio of the mean of the dependent element to the mean of the driver. As noted earlier, a limited number of probability distributions can be placed on the factor when there is uncertainty as to its magnitude.

The final input workscreen is Groupings. This is the vehicle by which an analyst can identify subsets of elements that move together, either positively or negatively, and the strengths of their relationships. It is especially important if certain correlations have been determined off-line either analytically or by simulation. However, we have pointed out what we consider to be limitations in the way RI\$K accepts and processes this information.

Having provided all inputs required, the final step is to call for the Calculation routine. As noted earlier, a user may choose between (or examine both) an analytic solution and a Monte Carlo simulation. If the simulation is selected, the user may specify the number of iterations (random drawings) desired. Both tabular and graphical output are available. Figure 2 is an example of the graphs of a probability distribution and cumulative distribution of total cost in a hypothetical uncertainty analysis.

9. Strictly speaking, factors can be applied without use of the Factor Specification screen. Recalling an example from the section of the paper on analytical preliminaries, where support cost was a fixed fraction of hardware cost, the only inputs that RI\$K requires for support costs are (1) its mean, dispersion measure, and distribution form, all of which can be easily determined from knowledge of the driver element, and (2) the information (via the Groupings workscreen) that hardware and support are correlated at exactly 1.0.

Figure 2. RI\$K graphical output



Crystal Ball and NCAP

Crystal Ball is a software package that supplements the capabilities of spreadsheets such as Excel and Lotus 1-2-3. It permits the user to define random variables in the spreadsheet and provides a limited Monte Carlo capability. This section provides a detailed description and evaluation of Crystal Ball by comparing its features with those of the cost uncertainty package in current use at the Naval Center for Cost Analysis—the package we refer to simply as NCAP.

NCAP was written in spreadsheet form using Lotus 1-2-3, version 3.1. Lotus macros perform the generation of random numbers, the Monte Carlo simulations, and the data analysis. We obtained it on a diskette along with instructions for its operation.

To provide a basis for comparison, we examined an actual cost uncertainty analysis using each of the packages. The original uncertainty analysis was carried out in connection with the Cooperative Engagement Capability (CEC) cost analysis performed by analysts at the cost center. We used a CNA personal computer (Gateway 2000, model P4D-33 with 486 processor) in examining each package. We first typed the data from the CEC analysis into a Lotus (version 3.1) spreadsheet. Then, after loading the Crystal Ball software into an Excel (version 5.0) spreadsheet, we copied the Lotus data into the Excel spreadsheet. (Our use of Excel should not be taken as an endorsement of that software over Lotus. Informally, we understand that the latest models of each product are similar in many respects.)

Detailed comparison of Crystal Ball and NCAP

This section compares the two packages with respect to (1) documentation, (2) running time, (3) size limitations, (4) number of variables analyzed, (5) ability to handle correlations, and (6) probability distributions of random variables.

Documentation

NCAP is almost completely undocumented. Its operation requires the uploading of information from a diskette or other computer file, and coaching from a knowledgeable user. Crystal Ball is a commercial product with a user's manual [4] and the usual technical support by telephone.

Another aspect of documentation is the ability to document and archive any given analysis. This involves documenting the input data, formulas used, and other facets of the analysis. Table 4 is a Crystal Ball output that provides a complete record of the CEC inputs, equations, distribution forms, etc., on a single page—including notes adequate to reproduce the results. NCAP would require storage of a diskette along with a page or pages of other information.

Running time

For the Crystal Ball example in table 4, the running time for 1,000 Monte Carlo iterations is about 45 seconds. Two thousand iterations require about 1 minute, 15 seconds. Informal estimates from NCAP users suggest its corresponding run time is 7 or 8 minutes for 1,000 iterations. Part of the difference in run time may be attributable to the use of an older version (3.1) of Lotus with NCAP, and the most recent version (5.0) of Excel with Crystal Ball. Another part of the difference may be attributable to Crystal Ball's being written in Turbo Pascal, whereas NCAP is in Lotus macro language.

Size limitations

NCAP is size-limited to one spreadsheet page, about 25 usable lines for instructions. Analyses requiring more than 25 lines must be broken into parts that can be run sequentially. Crystal Ball imposes no limitation on the number of lines.

Number of variables analyzed

NCAP provides a complete analysis on at most two variables, although partial information (sample means, standard deviations, and coefficients of variation) is provided for all variables. Crystal Ball permits as many variables as desired to be selected for analysis. In some respects, this is a trifling difference because the user is probably interested in only one or two variables. However, figure 3 illustrates how this feature might be of interest to the analyst.

Table 4. Example documentation from Crystal Ball

A	B	C	D	E	F	G	H	I	J	K	L	M
1												
2												
3	Nr	Mean	Std Dev	Value								
4	HW (DDS/CEP)	2677	0.001	=(K4*D4+C4)*6								
5	MINIT FACTOR	0.152	=0.15/05	=K5*D5+C5								
6	MINIT HW (DDS/CEP)			=(1-E5)*E4								
7	Antenna Enc. HW	383.98	0.001	=(K7*D7+C7)*6								
8	Tot HW to be MILSP			=E7+E6								
9	MILSPEC factor	1.51	0.541	=K9*D9+C9								
10	MILSPEC Adj.			=E9*E8*0.15								
11	HW w MILSPEC Adj.			=E10+E8								
12	TR Module T1(6chip)	24.862	0.001	=K12*D12+C12								
13	STEP-DOWN	0.79	0.134	=K13*D13+C13								
14	TR Mod T1 w ST-DN			=E12+E13								
15	MINIT Factor	0.152	0.023	=K15*D15+C15								
16	TR Module T1(3chip)			=(1-E15)*E14								
17	90.1%LC - B	-0.15	0.034	=K17*D17+C17								
18	TR Module HW			=282.6+E17+E18+E16								
19	ECU & BBU	521.9	78.285	=(K19*D19+C19)*5								
20	Total EMD HW			=E19+E18+E11								
21	Air EMC HW			=E11/6+(216/616)*E18								
22	Contractor Supp. CER	1.098	0.176	=K22*D22+C22								
23	Contractor Supp.			=E22*E20-490.115								
24	Spares Factor	0.113	0.011	=K24*D24+C24								
25	Spares			=E24*E20								
26	Design CER	0.64	0.18	=K26*D26+C26								
27	Design			=E26*E20-15.6								
28	Total to ECI less fees			=E27+E25+E23+E20								
29												
30												
31	G&A, COM,FEE Factor	0.259	0.026	=K31*D31+C31								
32	G&A, COM, FEE			=E31*E28								
33	Total to ECI			=E32+E28								
34	E-2C Software	52020	7283	=K34*D34+C34								
35	All Other Software	475428	71314	=K35*D35+C35								
36	Total EMD Cont. Costs			=E36+E34+E33								
37	GIH CER	0.109	0.034	=K37*D37+C37								
38	GIH			=E37*E36-3647								
39	CASS	15000	2250	=K39*D39+C39								
40	E-2C Inst. Factor	0.322	0.033	=K40*D40+C40								
41	E-2C Inst.			=E40*E21								
42	SHIP Inst.	8650	2716	=K42*D42+C42								
43	E-2C Integration	72905	54679	=K43*D43+C43								
44	Total EMD			=SUM(E41:E43)+E39+E38+E36								

Figure 3. Crystal Ball Report

Crystal Ball Report
Simulation started on 3/31/95 at 12:40:32
Simulation stopped on 3/31/95 at 12:41:22




Figure 3 provides a breakdown of Total EMD cost, a primary output variable in the CEC example. This variable is the sum of the six variables listed in the left column, and the figure displays the contribution of each of the six variables to the total variance of the Total EMD cost. This variance decomposition is similar in spirit to that of an analysis of variance (ANOVA) table. For example, figure 3 shows that over 60 percent of the variability in Total EMD Cost is due to the single input "EMD Cont. Costs." Also, the bottom three inputs (E2C Inst, CASS, and Ship Inst) contribute less than one percent to the variability of the Total EMD Cost. These latter three variables could probably be entered as simple constants in the spreadsheet without changing the results of the analysis. The user might then choose one of the six input variables for a further variance breakdown. Such analysis may provide an increased understanding of the underlying cost model.

Ability to handle correlations

In practice, cost variables are usually correlated due to some direct relationship between the variables or to underlying factors that are common to both. NCAP can create correlation between cost variables only by simulating a direct relationship or common factor as an explicit piece in the spreadsheet. A Crystal Ball model can also generate correlation between cost variables in this way, but in addition, the user may simply specify two variables as being correlated with a desired correlation coefficient, and Crystal Ball will simulate these correlated random variables without reference to the spreadsheet.

Usually, correlation is measured using the standard statistical "correlation coefficient." It is well known [5] that it may be difficult to generate random numbers having a desired joint distribution with a desired correlation matrix; in fact, it may be impossible unless there are appropriate bounds on the elements in the correlation matrix. Crystal Ball avoids some of these problems by using "rank correlation" techniques [6] which provide a rapid, nonparametric approach with a slight loss of efficiency. If the user specifies correlation values that are impossible (with the given marginal distributions), then Crystal Ball approximates the correlation values as closely as possible and

prints out a warning message.¹⁰ (The authors have not tested this feature.) The reader is reminded of the earlier discussion concerning (1) the controversy associated with use of subjective measures of correlation, and (2) the potential value of "what if" or sensitivity analysis in the area.

Distributions of random variables

The NCAP software provides the user with a choice of four distributions (normal, *t*, uniform, and custom). Modification of distribution parameters is by keyboard entry only. Crystal Ball provides 16 distributions. Modification of parameters is either by keyboard or graphically—using the mouse.

There is little agreement in the cost analysis literature as to what distributions should be used or how parameters should be selected. However, regardless of dispute, NCAP appears to be too limited in its offering of distributions. As noted previously, when the mean of a variable is close to zero or the standard distribution is similar in size to the mean, then a normal or *t* random variable will take on negative values a nontrivial fraction of the time. In a cost analysis context, negative costs are usually unrealistic, and it is desirable to have readily available a log-normal or some other distribution to accommodate these variables.

Concluding remarks

The foregoing comparison establishes a reasonably strong basis for selecting Crystal Ball over NCAP as a spreadsheet overlay. However, one additional point should not go unmentioned. NCAP is presently available to, and in fact is being used by, the cost center. Crystal Ball, while relatively inexpensive, must nevertheless be acquired through a formal procurement action.

10. When two random variables are related through a joint probability distribution, the *marginal* distribution of each is simply each variable's univariate distribution.

Appendix A: Literature review

Introduction

This appendix reports on the results of a brief literature search in the area of cost-risk analysis. The emphasis of the search was primarily to obtain a perspective on the state of the art in this area. A secondary purpose was to become acquainted with the capabilities of the cost-risk analysis models that appear to be in common use.

As a result of the search, several themes emerged:

- There is no generally accepted definition of “risk” in cost-risk analysis.
- No single methodology emerges as being “best.”
 - Monte Carlo methodology is considered to be one of the better approaches but is not a panacea.
 - Decision analysis is often touted as being the theoretically best approach, but the details of its practical use do not appear to be well known in the analysis community.
- There is general agreement that cost estimates must include information about their associated probability distributions. A variety of distributions are in common use, with the beta and triangular being the most common. There is no agreement on the best way to estimate parameters for these distributions.
- There is general agreement that correlations between costs must be considered in estimating total cost. There is no general agreement about how to obtain estimates of such correlations or how to incorporate this information into the analysis.

- There is some feeling that a critical (and often overlooked) part of performing a cost-risk assessment is obtaining appropriate input data.

Some frequently cited software packages or models are listed below. Each is described later in the appendix, and several of the packages are compared in [7].

- CLT Central Limit Theorem (USASSDC)
- FRISKEM Formal Risk Evaluation Methodology (Aerospace)
- PACER Parametric Cost-Estimating Relationship Module (DSMC)
- @RISK Spreadsheet Add-in Model (Palisade)
- Crystal Ball Spreadsheet Add-in Model (Decisioneering)
- RI\$K Cost Risk Model (Tecalote Research)

Discussion

Definition of "risk"

The cost-analysis community seems to have no generally accepted definition of "risk." Some writers appear to regard risk and uncertainty as synonymous, whereas others define risk rigidly in some statistical framework—as in decision theory [8, 9, 10]. One sees terms such as "cost," "risk," "cost risk," "cost-estimating risk," "project risk," "schedule risk," and "technical risk" used rather informally.

The underlying theme that the literature conveys is that there is always an implied "best estimate" and the associated "risk" is some measure of the extent to which the actual result may overshoot the "best estimate." "Cost risk analysis" is the process of generating this "best estimate" and associated "risk." Traditionally, the "best estimate" is most important and is calculated first, after which some estimate of "risk" is made.

The modern trend is to decry this split, and to argue that the uncertainty in a cost estimate is as important, if not more so, than any point estimate [9, 10]. All of the software packages examined in this literature search made some attempt to quantify cost uncertainty as well as

provide a point estimate. Percentiles of the total cost distribution are commonly used to help quantify cost uncertainty.

No methodology is "best"

The general approach to "cost risk analysis" seems to be a bottom-up approach. The analyst attempts to obtain or generate cost estimates for all of the individual cost elements, then sum these estimates into a total cost estimate, from which a point estimate and associated risk measures can be obtained. Each step gives rise to difficulties [8].

A simple approach is to obtain three estimates for each of the individual element costs: best case, worst case, best estimate. These are summed over the individual elements to obtain these three estimates for the total cost. The worst case estimate is obtained from the best estimate by multiplying by some experience-based factor. This approach has the advantage of giving rapid results, which may however be difficult to defend.

A more sophisticated approach is to obtain or generate distributional information about each of the individual elements. These distributions are then combined to provide an estimate of the resulting total cost distribution. The combining of these distributions is a major problem with this approach. If the costs are independent (which they never are), then the distributions can be combined via repeated convolutions. But this is hard to do in practice because of computational difficulties. However, the means and variances of the individual cost elements can be summed to provide estimates of the total cost mean and variance. When individual element costs are correlated, this does not affect the estimate of the total cost mean, but the variance computation must include the correlation terms. Unfortunately, this gives only the first two moments of the total cost distribution. Alternatively, the individual element distributions can be combined via Monte Carlo into a total cost distribution. This can involve lengthy computation and resulting sampling errors. Also, it can be difficult to generate appropriately correlated Monte Carlo variates.

Some writers argue passionately that decision theory provides the only justifiable framework for quantification of total cost and definition of associated risks [9, 10, 11]. However, the methodology does

not appear to have filtered down from the classroom to the practitioner in any generally accepted form as yet.

Probability distributions

The modern view is that each cost estimate must include information about its probability distribution. When there are ample data, then statistical and curve-fitting techniques exist for determining which probability distribution provides a best fit. At worst, an empirical distribution can be generated. But generally, such ample data do not exist, and the analyst must make do with scraps of distributional information, including expert judgment. There is a great deal of discussion in the literature about the benefits of using the triangular and beta distributions. Because of their finite upper and lower values and modal behavior, they can be fitted using only the three estimates: worst case, best case, and best estimate. There are some possible problems in finding the best fit, but the method is popular and is mechanized in some of the computer models that were examined. Other distributions are in common use, with no general agreement in the analytic community about which distributions are best for what purposes [8, 12]. Most of the models provide for automated fitting with a variety of distributions.

Additional problems arise when expert judgment is the basis for estimation of a probability distribution. It appears to be well established that human beings are not very good at estimating probabilities, particularly tail probabilities, because of numerous biases that seem to be "wired in." Even well trained statisticians have these biases. As a result, the upper and lower tails in probability distributions are almost always underestimated, and associated "risk" is almost always underestimated, no matter how it is defined [8].

When Bayesian methods are used, it may be necessary to encode subjective probability distributions. This has been an area of intensive research, and some authors feel that current methodology may be far behind the state of the art [10].

Correlations between costs

There is general agreement that correlations between costs must be included in the analysis. There are two problems: How do you get the correlation estimates, and how do you include them? There are few examples of correlation estimates in the literature, but there is much discussion of how certain processes should be correlated with other processes. Including the correlation estimates in the analysis can be difficult when Monte Carlo methods are involved. It can also be difficult to simulate random numbers with precisely the right marginal distributions and precisely the right correlations, except in a few special cases. This appears to be an active area of current research [5, 13 through 16].

Data collection and analysis

Many of the technical problems discussed above are exacerbated by the difficulty of getting reliable data. Some authors feel that this is a neglected area that is crucial [8]. Reference [17] illustrates the effort that is needed to collect, organize, and sanitize large bodies of data.

Comparison of selected software packages

CLT—Central Limit Theorem (USASSDC) [7]

General approach: If the detailed cost elements are independent with finite means and variances, then the sum of the means and the sum of the variances are equal to the arithmetic mean and variance of the total system cost. Also, the distribution of the total cost approaches that of a normal distribution as the number of detailed cost elements increases.

Implementation: Written in BASIC for the PC.

Built-in distributions: Beta, triangular, uniform, normal.

Strengths: Analytic model, fast-running.

Weaknesses: Does not allow correlation between cost elements. Not compatible with any spreadsheet or work processor software.

FRISKEM—Formal Risk Evaluation Methodology (Aerospace) [18]

General approach: For each WBS element, FRISKEM accepts low, best, and high cost estimates, along with an interelement correlation matrix. The model fits a triangular distribution of cost to each WBS element and calculates the mean and variance of each triangular distribution. Using these with the correlation matrix, the mean and variance of the total cost distribution are obtained. These parameters determine a lognormal distribution, from which cost-risk measures can be obtained. This model extends an earlier model FRISK.

Built-in distributions: Triangular, lognormal.

Implementation: Written in BASIC for the PC.

Strengths: Analytic model. Fast-running. Does allow correlation between costs. Easy to learn.

Weaknesses: Uses only triangular distributions to fit the WBS cost elements. The total cost distribution is hardwired to be lognormal.

PACER—Parametric Cost-Estimating Relationship Module (DSMC) [7]

General approach: PACER is a “tool box” of standalone applications, with four subsystems: Utility, Cost-Estimating Relationships, Operating, and Applications. The risk analysis function is a subroutine of the Applications System and is based on the Central Limit Theorem. (See CLT description below.)

Built-in distributions: Six precalculated beta distributions.

Implementation: Written in C for the PC.

Strengths: Analytic model. Fast-running. Compatible with some spreadsheet and word processing software.

Weaknesses: Does not allow correlation between costs. Not easy to learn.

@RISK—Spreadsheet Add-in Model (Palisade Corp.) [7]

General approach: @RISK is a simulation model that uses either Monte Carlo or Latin Hypercube sampling.

Built-in distributions: Over 30 types of distributions.

Implementation: Can be added to either Excel or Lotus on PC or Mac.

Strengths: Very flexible. Good tabular and statistical outputs. Cost elements may be correlated.

Weaknesses: The model assumes a sophisticated user. Slow execution time.

Crystal Ball—Spreadsheet Add-in Model (Decisioneering, Inc.) [4, 19]

General approach: Crystal Ball is a simulation model that uses either Monte Carlo or Latin Hypercube sampling.

Built-in distributions: Sixteen types of distributions.

Implementation: Can be added to either Excel or Lotus on PC or Mac.

Strengths: Very flexible. Good tabular and statistical outputs. Easy to learn. Permits correlated cost elements.

Weaknesses: Moderately slow execute time.

RI\$K—Cost Risk Model (Tecolote Research) [7]

General approach: RI\$K can be run either as a Monte Carlo simulation model or as an analytic model. The analytic model assumes that the total cost distribution can be modeled as beta.

Built-in distributions: Normal, lognormal, beta, triangular, uniform.

Implementation: Written in C for the PC (Windows compatible).

Strengths and weaknesses: See the section on RI\$K in the main body of the paper.

Appendix B: Parameter computations involving the product of two random variables

In the simple example of cost uncertainty analysis presented in the main body of the paper, one case that was considered had support cost (S) estimated as a factor of hardware cost (H), with the factor (F) assumed to be a uniformly distributed random variable. Hence, $S = FH$. Here we illustrate how the variance and standard deviation of S , the correlation between S and H , and the standard deviation of total cost (TC) may be computed analytically. We begin by reviewing certain definitions and properties of random variables.

The mean, or expected value, of a random variable X is denoted by

$$E(X) = \mu .$$

The variance of X is given by

$$E(X - \mu)^2 = E(X^2) - \mu^2 = \sigma^2 ,$$

with the standard deviation being simply σ . For the product of two *independent* random variables, X_1 and X_2 , the mean, variance, and covariance are defined as follows:

$$E(X_1 X_2) = E(X_1) E(X_2) = \mu_1 \mu_2$$

$$Var(X_1 X_2) = E[X_1 X_2 - E(X_1 X_2)]^2 = E(X_1^2) E(X_2^2) - \mu_1^2 \mu_2^2$$

$$Cov(X_1 X_2) = E(X_1 - \mu_1) (X_2 - \mu_2) = E(X_1 X_2) - \mu_1 \mu_2 .$$

Note that, because X_1 and X_2 are independent, their covariance is identically zero. If the variables were not independent, $E(X_1 X_2) \neq \mu_1 \mu_2$, and the covariance would be either positive or

negative. The standardized covariance (correlation) between any two random variables, X_1 and X_2 , denoted by ρ_{12} , is given by

$$\rho_{12} = \text{Cov}(X_1, X_2) / \sigma_1 \sigma_2 .$$

In the cost uncertainty example, the mean and standard deviation of H , μ_H and σ_H , were assigned values of 100 and 20, respectively. It therefore follows that

$$E(H^2) = \sigma_H^2 + \mu_H^2 = 10,400 .$$

The variable factor F was assumed to be independent of H and uniformly distributed over the interval $[0.35, 0.65]$. For any uniform variable U distributed over the interval $[a, b]$, properties of relevance here—as developed in, for example, [20, pp. 297-298]—are

$$E(U) = (a + b) / 2$$

$$E(U^2) = [(a + b)^2 - ab] / 3$$

$$\sigma_U^2 = (b - a)^2 / 12 .$$

Applying these results to the uniformly distributed factor F with $a = 0.35$ and $b = 0.65$, we obtain

$$E(F) = 0.5, \quad E(F^2) = 0.2575, \quad \sigma_F^2 = 0.0075 .$$

Given the above, and recalling that $S = FH$ with F and H assumed to be independent, we may compute the following:

$$E(S) = E(FH) = E(F)E(H) = \mu_S = 50$$

$$E(S^2) = E(FH)^2 = E(F^2)E(H^2) = 2,678$$

$$\sigma_S^2 = E(S^2) - \mu_S^2 = 178$$

$$\sigma_S = (178)^{1/2} = 13.34 .$$

The covariance of H and S is given by

$$\begin{aligned} Cov(H, S) &= E(HS) - \mu_H \mu_S \\ &= E(H^2 F) - \mu_H \mu_S \\ &= E(H^2) E(F) - \mu_H \mu_S \\ &= 200 . \end{aligned}$$

The correlation between H and S is therefore

$$\rho_{HS} = Cov(H, S) / \sigma_H \sigma_S = 0.75 .$$

Finally, as noted in the main text, the standard deviation of total cost is given by

$$\sigma_{TC} = (\sigma_H^2 + \sigma_S^2 + 2\rho_{HS}\sigma_H\sigma_S)^{1/2} = (978)^{1/2} = 31.27 .$$

References

- [1] Air Force Institute of Technology Paper No. AU-AFIT/LA-TR-94-1, *An Investigation of the Accuracy of Heuristic Methods for Cost Uncertainty Analysis*, by Wendell P. Simpson and Kevin P. Grans, Aug 1994
- [2] CNA Research Memorandum 93-228, *A Study of the Predictive Accuracy of Alternatively Estimated Statistical Cost Models*, by Henry L. Eskew and Kletus S. Lawler, Dec 1993
- [3] Naval Center for Cost Analysis, *Standard Cost Factors Handbook*, Nov 1992
- [4] Decisioneering, *Crystal Ball Version 3.0 User Manual*, 1993
- [5] Institute for Defense Analysis Paper P-2998, *A Method for Simulating Correlated Random Variables from Partially-Specified Distributions*, by P. M. Lurie and M. S. Goldberg, Oct 1994
- [6] R. L. Iman and W. J. Conover. "A Distribution-Free Approach to Inducing Rank Correlation Among Input Variables." *Communications in Statistics: Simulation and Computation*, Vol. 2, No. 3, 311-334, 1982
- [7] Tecolote Research, Inc., Report CR-0644, *Risk Model Comparisons for Cost Analysis*, by W. H. Jago, M. K. Allen, and L. S. Fichter, May 1993
- [8] F. Biery, D. Hudak, and Lansdowne. "Improving Cost Risk Analyses." *Journal of Cost Analysis*, Spring, 1994
- [9] J. H. Schuyler. "Credible Projections Now Require Decision Analysis." *Cost Engineering*, Vol. 34, No. 3, Mar 1992

- [10] K. T. Wallenius. *Cost Uncertainty Assessment Methodology: A Critical Overview*, Technical Report 491, Dept of Mathematical Sciences, Clemson University, Aug 1985
- [11] D. Samson. *Managerial Decision Analysis*. Irwin, 1988
- [12] K. T. Wallenius, *Cost Uncertainty Assessment Methodology: New Initiatives*, Technical Report 492, Dept of Mathematical Sciences, Clemson University, Aug 1985
- [13] S. A. Book and P. H. Young. *Monte-Carlo Generation of Total-cost Distributions When WBS Element Costs Are Correlated*, submitted 24th annual DOD Cost Symposium, 1990
- [14] P. R. Garvey and A. E. Taub. *A Joint Probability Model for Cost and Schedule Uncertainties*, The Mitre Corporation, presented at the 26th annual DOD Cost Analysis Symposium, Sep 1992
- [15] IDA Paper P-2732, *Simulated Correlated Distributions With Bounded Domains*, by P. H. Lurie and M. S. Goldberg, Sep 1992
- [16] P. R. Garvey. "A General Analytic Approach to System Cost Uncertainty Analysis." *Cost Analysis and Estimating*, Eds. W. R. Greer and D. A. Nussbaum, Springer-Verlag, 1990
- [17] R. A. Katz. "Parametric CERs for Replenishment Repair Parts." *Cost Analysis and Estimating*, Eds. W. R. Greer and D. A. Nussbaum, Springer-Verlag, 1990
- [18] R. L. Abramson and P. H. Young. *FRISKEM—Formal Risk Evaluation Methodology*, submitted to *Journal of Cost Analysis*
- [19] Crystal Ball, (Software review) *APICS—The Performance Advantage*, Jan 1994
- [20] R. C. Pfaffenberger and J. H. Patterson. *Statistical Methods*. Irwin, 1987

Distribution list

SNDL

A1 DASN - AIR PROGRAMS
A1A DASN - MANPOWER
A1B UNSECNAV
A1F ASSTSECNAV FM WASHINGTON
Attn: NCA
A1J ASSTSECNAV RDA WASHINGTON
A1J DASN - C4I/EW/SPACE
PROGRAMS
A1J DASN - MUW
A1J DASN - SHIP PROGRAMS
A2A ONR
A2A OPA
A6 HQMC P&R (2 COPIES)
FF38 USNA ANNAPOLIS
FF42 NAVPGSCOL MONTEREY CA
FKA1A COMNAVAIRSYSCOM DC
Attn: AIR 4.2
FKA1B COMSPAWARSYSCOM DC
FKA1C COMNAVFACEGCOM VA
FKA1F COMNAVSUPSYSCOM DC
FKA1G COMNAVSEASYSYSCOM DC
Attn: SEA-017
FKP4D NAVSURFWARCEN
CARDEROCKDIV BETHESDA

FKP4E NAVSURFWARCENDIV
FKR6A NAVAIRWARCENACDIV
V12 CG MCCDC STUDIES AND
ANALYSES DIVISION

OPNAV

N4B
N801
N801C
N810
N811
N812
N812D
N863
N880B

OTHER

IDA
RAND SANTA MONICA
LMI
ARMY CEAC
AIR FORCE CAA